

Counting Unique-Sink Orientations

Jan Foniok^{*†}

Bernd Gärtner^{‡§}

Lorenz Klaus^{*¶}

Markus Sprecher^{‡||}

7th December 2010

Abstract

Unique-sink orientations (USOs) are an abstract class of orientations of the n -cube graph. We consider some classes of USOs that are of interest in connection with the linear complementarity problem. We summarise old and show new lower and upper bounds on the sizes of some such classes. Furthermore, we provide a characterisation of K-matrices in terms of their corresponding USOs.

Keywords: unique-sink orientation, linear complementarity problem, pivoting, P-matrix, K-matrix

MSC2010: 90C33, 52B12, 05A16

1 Introduction

Unique-sink orientations (USOs) are an abstract class of orientations of the n -cube graph. A number of concrete geometric optimisation problems can be shown to have the combinatorial structure of a USO. Examples are the linear programming problem [10], and the problem of finding the smallest enclosing ball of a set of points [10, 25], or a set of balls [6]. In this paper, we count the USOs of the n -cube that are generated by P-matrix linear complementarity problems (P-USOs). This class covers many of the “geometric” USOs. We show that the number of P-USOs is $2^{\Theta(n^3)}$. The lower bound construction is the interesting contribution here, and it even yields USOs from the subclass of K-USOs whose combinatorial structure is known to be very rigid [7]. In contrast, the number of all n -cube USOs is doubly exponential in n [13].

^{*}Institute for Operations Research, ETH Zurich, 8092 Zurich, Switzerland

[†]foniok@math.ethz.ch

[‡]Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland

[§]gaertner@inf.ethz.ch

[¶]lklaus@ifor.math.ethz.ch

^{||}markussp@student.ethz.ch

Unique-sink orientations

We follow the notation of [7]. Let $[n] := \{1, 2, \dots, n\}$. For a bit vector $v \in \{0, 1\}^n$ and $I \subseteq [n]$, let $v \oplus I$ be the element of $\{0, 1\}^n$ defined by

$$(v \oplus I)_j := \begin{cases} 1 - v_j & \text{if } j \in I, \\ v_j & \text{if } j \notin I. \end{cases}$$

Instead of $v \oplus \{i\}$ we write $v \oplus i$.

Under this notation, the (undirected) n -cube is the graph $G = (V, E)$ with

$$V := \{0, 1\}^n, \quad E := \{\{v, v \oplus i\} : v \in V, i \in [n]\}.$$

A *subcube* of G is a subgraph $G' = (V', E')$ of G where $V' = \{v \oplus I : I \subseteq C\}$ for some vertex v and some set $C \subseteq [n]$, and $E' = E \cap \binom{V'}{2}$. The *dimension* of such a subcube is $|C|$.

Let ϕ be an orientation of the n -cube (a digraph with underlying undirected graph G). If ϕ contains the directed edge $(v, v \oplus i)$, we write $v \xrightarrow{\phi} v \oplus i$, or simply $v \rightarrow v \oplus i$ if ϕ is clear from the context. If V' is the vertex set of a subcube, then the directed subgraph of ϕ induced by V' is denoted by $\phi[V']$. For $F \subseteq [n]$, let $\phi^{(F)}$ be the orientation of the n -cube obtained by reversing all edges in coordinates contained in F ; formally

$$v \xrightarrow{\phi^{(F)}} v \oplus i \quad :\Leftrightarrow \quad \begin{cases} v \xrightarrow{\phi} v \oplus i & \text{if } i \notin F, \\ v \oplus i \xrightarrow{\phi} v & \text{if } i \in F. \end{cases}$$

An orientation ϕ of the n -cube is a *unique-sink orientation* (*USO*) if every subcube $G' = (V', E')$ has a unique sink (that is, vertex of outdegree zero) in $\phi[V']$. It is not difficult to show that in a unique-sink orientation, every subcube also has a unique source (that is, vertex of indegree zero).

A special USO is the *uniform orientation*, in which $v \rightarrow v \oplus i$ if and only if $v_i = 0$.

Unique-sink orientations enable a graph-theoretic description of simple principal pivoting algorithms for linear complementarity problems. They were introduced by Stickney and Watson [24] and have recently received much attention [9, 10, 13, 17, 22, 23, 25].

Linear complementarity problems

A *linear complementarity problem* ($\text{LCP}(M, q)$) is for a given matrix $M \in \mathbb{R}^{n \times n}$ and a vector $q \in \mathbb{R}^n$, to find vectors $w, z \in \mathbb{R}^n$ such that

$$w - Mz = q, \quad w, z \geq 0, \quad w^T z = 0. \quad (1)$$

A *P-matrix* is a square real matrix whose principal minors are all positive. If M is a P-matrix, the appertaining LCPs are called *P-LCPs*; in this case there exists a unique solution for any q .

Let $B \subseteq \{1, 2, \dots, n\}$, and let A_B be the $n \times n$ matrix whose i th column is the i th column of $-M$ if $i \in B$, and the i th column of the $n \times n$ identity matrix I_n otherwise. If M is a P-matrix, then A_B is invertible for every set B . We call B a *basis*. If $A_B^{-1}q \geq 0$, let

$$w_i := \begin{cases} 0 & \text{if } i \in B \\ (A_B^{-1}q)_i & \text{if } i \notin B \end{cases}, \quad z_i := \begin{cases} (A_B^{-1}q)_i & \text{if } i \in B \\ 0 & \text{if } i \notin B \end{cases}. \quad (2)$$

The vectors w, z are then a solution to the LCP (1).

A problem $\text{P-LCP}(M, q)$ is *nondegenerate* if $(A_B^{-1}q)_i \neq 0$ for all B and i . Following [24], a nondegenerate $\text{P-LCP}(M, q)$ induces a USO: For $v \in \{0, 1\}^n$, let $B(v) := \{j \in [n] : v_j = 1\}$. Then the unique-sink orientation ϕ induced by $\text{P-LCP}(M, q)$ is given by

$$v \xrightarrow{\phi} v \oplus i \quad :\Leftrightarrow \quad (A_{B(v)}^{-1}q)_i < 0. \quad (3)$$

The run of a simple principal pivoting method (see [19, Chapter 4]) for the P-LCP then corresponds to following a directed path in the orientation ϕ . Finding the sink of the orientation is equivalent to finding a basis B with $A_B^{-1}q \geq 0$, and thus via (2) to finding the solution to the P-LCP.

In this paper, we are primarily interested in establishing bounds for the number of n -dimensional USOs satisfying some additional properties (for instance, USOs induced by P-LCPs), which we introduce in the next section.

2 Matrix classes and USO classes

It is NP-complete to decide whether a solution to an LCP exists [2]. If the matrix M is a P-matrix, however, a solution always exists. The problem of finding it is unlikely to be NP-hard, because if it were, then $\text{NP} = \text{co-NP}$ [14]. Even so, no polynomial-time algorithms for solving P-LCPs are known. Hence our motivation to study some special matrix classes and investigate what combinatorial properties their USOs have. The ultimate goal is then to try and exploit these combinatorial properties in order to find an efficient algorithm for the corresponding LCPs.

A *Z-matrix* is a square matrix whose all off-diagonal elements are non-positive. A *K-matrix* is a matrix which is both a Z-matrix and a P-matrix. A *hidden-K-matrix* is a P-matrix such that there exist Z-matrices X and Y and non-negative vectors r and s with $MX = Y$, $r^T X + s^T Y > 0$.

The importance of these matrix classes is due to the fact that polynomial-time algorithms are known for solving the $\text{LCP}(M, q)$ if the matrix M is a Z-matrix [1, 21], a hidden-K-matrix [12], or the transpose of a hidden-K-matrix [20].

A USO is a *P-USO* if it is induced via (3) by some $\text{LCP}(M, q)$ with a P-matrix M ; it is a *K-USO* if it is induced by some $\text{LCP}(M, q)$ with a K-matrix M ; and it is a *hidden-K-USO* if it is induced by some $\text{LCP}(M, q)$ with a hidden-K-matrix M .

A USO is a *Holt-Klee USO* if in each of its subcubes, there are d directed paths from the source to the sink of the subcube, with no two paths sharing a vertex other than

source and sink; here d is the dimension of the subcube. A USO ϕ is *strongly Holt–Klee* if $\phi^{(F)}$ is Holt–Klee for every $F \subseteq [n]$. By [9], every P-USO is a strongly Holt–Klee USO.

Finally, a USO is *locally uniform*, if

$$\begin{aligned} \text{whenever } u_i = u_j = 0 \text{ and } u \xrightarrow{\phi} u \oplus i, \ u \xrightarrow{\phi} u \oplus j, \\ \text{then } u \oplus i \xrightarrow{\phi} u \oplus \{i, j\}, \ u \oplus j \xrightarrow{\phi} u \oplus \{i, j\} \end{aligned} \quad (4)$$

and

$$\begin{aligned} \text{whenever } u_i = u_j = 0 \text{ and } u \oplus i \xrightarrow{\phi} u, \ u \oplus j \xrightarrow{\phi} u, \\ \text{then } u \oplus \{i, j\} \xrightarrow{\phi} u \oplus i, \ u \oplus \{i, j\} \xrightarrow{\phi} u \oplus j. \end{aligned} \quad (5)$$

By [7], every K-USO is locally uniform, and every locally uniform USO is acyclic. Thus we have:

$$\begin{aligned} \text{K-USOs} \subseteq \text{locally uniform P-USOs} \subset \text{acyclic P-USOs} \subset \text{P-USOs} \\ \subset \text{strongly Holt–Klee USOs} \subset \text{Holt–Klee USOs}. \end{aligned}$$

The first inclusion is not known to be strict; see also Section 4.

An *LP-USO* is an orientation of the n -cube admitting a realisation as a polytope in the n -dimensional Euclidean space, combinatorially equivalent to the n -cube, such that there exists a linear function f and

$$v \xrightarrow{\phi} v \oplus i \quad \text{if and only if} \quad f(v \oplus i) > f(v).$$

It follows from [11, 16, 18, 20] that LP-USOs are exactly hidden-K-USOs, and we have:

$$\text{K-USOs} \subset \text{LP-USOs} = \text{hidden-K-USOs} \subset \text{acyclic P-USOs}.$$

In the next section we examine the numbers of n -USOs in the respective classes.

It is also possible to obtain USOs from completely general linear programs. The reduction in [10] yields *PD-USOs*, i.e., USOs generated by LCPs with symmetric positive definite matrices M . Since these are exactly the symmetric P-matrices [3, Section 3.3], we also have PD-USOs \subset P-USOs. The USOs that are obtained from the problem of finding the smallest enclosing ball of a set of points [8, Section 3.2] are “almost” PD-USOs in the sense that every subcube not containing the origin 0 is oriented by a PD-USO [15]. For the USOs from smallest enclosing balls of *balls* [6], we are not aware of a similar result.

3 Counting USOs

First counting results about USOs were obtained by Matoušek [13], who gave asymptotic bounds on the number of all USOs and acyclic USOs.

Next, Develin [4]—in order to show that the Holt–Klee condition does not characterise LP-USOs—proved that the number of n -dimensional LP-USOs is bounded from above by $2^{O(n^3)}$, whereas the number of Holt–Klee USOs is bounded from below by $2^{\Omega(2^n/\sqrt{n})}$.

Using similar means, we prove an upper bound of $2^{O(n^3)}$ on the number of P-USOs, and observe that a slight modification of Develin’s construction yields a lower bound of $2^{\binom{n-1}{\lfloor (n-1)/2 \rfloor}}$ for strongly Holt–Klee locally uniform USOs. Furthermore, we provide a construction of $2^{\Omega(n^3)}$ K-USOs. These results imply that the number of K-USOs, LP-USOs, as well as P-USOs, is $2^{\Theta(n^3)}$.

Previously known and new bounds on the number of n -dimensional USOs in the classes defined in the previous section are summarised in the following table:

class	lower bound	upper bound
K-USOs	$2^{\Omega(n^3)}$	
LP-USOs		$2^{O(n^3)}$ [4]
P-USOs		$2^{O(n^3)}$
strongly Holt–Klee USOs	$2^{\Omega(2^n/\sqrt{n})}$	
Holt–Klee USOs [4]	$2^{\Omega(2^n/\sqrt{n})}$	
locally uniform USOs	$2^{\Omega(2^n/\sqrt{n})}$	
acyclic USOs [13]	$2^{2^{n-1}}$	$(n+1)^{2^n}$
all USOs [13]	$n^{\Omega(2^n)}$	$n^{O(2^n)}$

3.1 An upper bound for P-USOs

Every P-USO is determined by the sequence $\sigma(M, q) = (\text{sgn}(A_{B(v)}^{-1}q)_i : v \in \{0, 1\}^n, i \in [n])$, which is a function of the P-matrix M and the right-hand side q . Furthermore, we are interested only in nondegenerate right-hand sides q , which means we are interested only in sequences containing no 0.

In order to be able to apply algebraic tools, we first derive an equivalent description of the USO, using only polynomials.

3.1 Lemma. *Each entry of the vector $\sigma(M, q)$ is the sign of a polynomial in the entries of M and q of degree at most n .*

Proof. The entries of the matrix $A_{B(v)}^{-1}$ can be computed as

$$(A_{B(v)}^{-1})_{rs} = \frac{1}{\det A_{B(v)}} (-1)^{r+s} M_{sr},$$

where M_{rs} is the determinant of the submatrix of $A_{B(v)}$ obtained by deleting the r th row and the s th column, which is a polynomial of degree at most $n-1$. Hence

$$(A_{B(v)}^{-1}q)_i = \frac{1}{\det A_{B(v)}} \sum_{s=1}^n q_s \cdot (-1)^{i+s} \cdot M_{si}.$$

Recall that $A_{B(v)}$ has $|B(v)|$ columns of $-M$ and $n - |B(v)|$ columns of the identity matrix; thus $\text{sgn det } A_{B(v)} = (-1)^{|B(v)|}$, since M is a P-matrix. Therefore

$$\text{sgn}(A_{B(v)}^{-1}q)_i = \text{sgn}\left((-1)^{|B(v)|} \cdot \sum_{s=1}^n q_s \cdot (-1)^{i+s} \cdot M_{si}\right),$$

which is the sign of a polynomial of degree at most n . \square

The algebraic tool we will apply is the following theorem.

3.2 Theorem (Warren [26]). *Let p_1, \dots, p_s be real polynomials in k variables, each of degree at most d . If $s \geq k$, then the number of sign sequences $\sigma(x) = (\text{sgn } p_1(x), \dots, \text{sgn } p_m(x))$ that consist of terms $+1, -1$ is at most $(4eds/k)^k$.*

Now all is set to prove an upper bound on the number of P-USOs.

3.3 Theorem. *The number of distinct n -dimensional P-USOs is at most $2^{O(n^3)}$.*

Proof. By Lemma 3.1, each P-USO is determined by a vector of $n2^n$ nonzero signs of polynomials of degree at most n . The number of variables is $n^2 + n$ (equal to the number of entries of the matrix M and the vector q). By Theorem 3.2, there are at most

$$\left(\frac{4e \cdot n \cdot n2^n}{n^2 + n}\right)^{n^2+n} \leq (4e \cdot 2^n)^{n^2+n} = 2^{O(n^3)}$$

such sign vectors. \square

3.2 A lower bound for strongly Holt-Klee and locally uniform USOs

A *monotone Boolean function* is a function $f : \{0, 1\}^k \rightarrow \{0, 1\}$ such that if $x \leq y$, then $f(x) \leq f(y)$. Counting monotone Boolean functions is known as *Dedekind's problem*. Let M be the set of $0, 1$ -vectors of length k with exactly $\lfloor k/2 \rfloor$ ones. A lower bound of $2^{\binom{k}{\lfloor k/2 \rfloor}}$ on the number of k -variate monotone Boolean functions can be obtained by taking for each subset $A \subseteq M$ the function f_A given by

$$f_A(x) = 1 \quad \text{iff} \quad \{y \in A : y \leq x\} \neq \emptyset.$$

3.4 Theorem. *The number of acyclic locally uniform strongly Holt-Klee n -USOs is at least $2^{\binom{n-1}{\lfloor (n-1)/2 \rfloor}} = 2^{\Omega(2^n/\sqrt{n})}$.*

Proof. Given an $(n-1)$ -variate monotone Boolean function f , we construct an n -USO ϕ by setting

$$\begin{aligned} v &\xrightarrow{\phi} v \oplus i \text{ if } i \neq n \text{ and } v_i = 0, \\ v &\xrightarrow{\phi} v \oplus n \text{ if } v_n + f(v') = 1, \end{aligned}$$

where $v' \in \{0, 1\}^{n-1}$ is formed by the initial $n - 1$ bits of v .

The orientation ϕ is clearly acyclic because any directed walk in ϕ is monotone on the first $n - 1$ bits. It is easy to show local uniformity too. The assumption of (5) is never satisfied. For (4) it suffices to consider the case $j = n$: If $u \xrightarrow{\phi} v \oplus n$, then $f(v') = 1$, hence $v \oplus i \xrightarrow{\phi} v \oplus \{i, n\}$.

For the strong Holt–Klee property, let $F \subseteq [n]$ and let $V' = \{v \oplus i : i \in C\}$ be the vertex set of a subcube with $|C| = d$. If $n \notin C$, then $\phi^{(F)}[V']$ is isomorphic to the uniform orientation, which is well-known to satisfy the Holt–Klee property. So suppose $n \in C$. Let $V_0 := \{v \in V' : v_n = 0\}$ and $V_1 := \{v \in V' : v_n = 1\}$ and let s be the source and t the sink of $\phi^{(F)}[V']$. Note that $\phi^{(F)}[V_0]$ and $\phi^{(F)}[V_1]$ are identical if we truncate the last coordinate of their vertices, and isomorphic to the uniform USO.

Now we distinguish two cases. First, if $b := s_n = t_n$, there are $d - 1$ disjoint paths from s to t in $\phi^{(F)}[V_b]$ and another path obtained by concatenating the edge $s \rightarrow s \oplus n$, a path in $\phi^{(F)}[V_{1-b}]$ from $s \oplus n$ to $t \oplus n$, and the edge $t \oplus n \rightarrow t$.

Second, let $b := s_n = 1 - t_n$. Without loss of generality we may assume that $b = 0$ and $n \notin F$. Let $P(i_1, \dots, i_d)$ denote the directed path $s \rightarrow s \oplus \{i_1\} \rightarrow s \oplus \{i_1, i_2\} \rightarrow \dots \rightarrow s \oplus \{i_1, i_2, \dots, i_d\}$. Order the elements of $C \setminus \{n\} = \{j_1, j_2, \dots, j_{d-1}\}$ so that for $j_k \in F$ and $j_\ell \notin F$ we have $k < \ell$. Since $\phi^{(F)}[V_0]$ and $\phi^{(F)}[V_1]$ are both isomorphic to the uniform orientation and $s_n \neq t_n$, we have $t = s \oplus C$. Now we claim that the paths $P(j_1, j_2, \dots, j_{d-1}, n)$, $P(j_2, j_3, \dots, j_{d-1}, n, j_1)$, \dots , $P(j_{d-1}, n, j_1, j_2, \dots, j_{d-2})$, $P(n, j_1, j_2, \dots, j_{d-1})$ are vertex-disjoint directed paths from s to t . The only non-obvious fact to show is that for any k , there is a directed edge $u := s \oplus \{j_k, j_{k+1}, \dots, j_{d-1}\} \rightarrow v := s \oplus \{j_k, j_{k+1}, \dots, j_{d-1}, n\}$. Note that $u \rightarrow v$ if and only if $f(u') = 1$ and that $f(s') = f(t') = 1$. If $j_k \notin F$, then $s' \leq u'$ and so $1 = f(s') \leq f(u')$, thus $f(u') = 1$. If on the other hand $j_k \in F$, then $t' \leq u'$ and so $1 = f(t') \leq f(u')$, thus $f(u') = 1$. Hence $u \rightarrow v$.

Therefore the number of acyclic locally uniform strongly Holt–Klee n -USOs is lower bounded by the number of $(n - 1)$ -variate monotone Boolean functions, which concludes the proof. \square

Remark. After swapping the roles of 0 and 1 in the n th coordinate, the above construction is the same as Mike Develin’s construction [4] of many orientations satisfying the Holt–Klee condition. Thus, both Develin’s and our construction yield Holt–Klee orientations, but local uniformity is obtained only in our variant.

The logarithm of the total number of acyclic n -USOs is no more than $2^n \log(n + 1)$ [13]. In comparison, the exponent in the lower bound obtained from Theorem 3.4 is of the order $2^n / \sqrt{n}$, and therefore still exponential. Restricting to K-USOs, the exponent goes down to a polynomial in n .

3.3 A lower bound for K-USOs

3.5 Theorem. *The number of distinct K-USOs in dimension n is at least $2^{\Omega(n^3)}$.*

Proof. Consider the upper triangular matrix

$$M(\beta) = \begin{pmatrix} 1 & -1 - \beta_{1,2} & -1 - \beta_{1,3} & \dots & -1 - \beta_{1,n} \\ 0 & 1 & -1 - \beta_{2,3} & \dots & -1 - \beta_{2,n} \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & -1 - \beta_{n-1,n} \\ 0 & 0 & 0 & \dots & 1 \end{pmatrix}$$

and the vector $q = (-1, 1, -1, \dots, (-1)^n)^T$. We will now examine how the choice of the parameters $\beta_{i,j}$ influences the USO induced by the LCP($M(\beta), q$). Our goal is to show that we can make $2^{\Omega(n^3)}$ choices, each of which induces a different USO.

First, let $B \subseteq [n]$ and, analogously to the definition of A_B , let $A_B(\beta)$ be the matrix whose i th column is the i th column of $-M(\beta)$ if $i \in B$, and the i th column of I_n otherwise. The reader is kindly invited to verify, by straightforward computation, that

$$\sigma_r \cdot ((A_B(\beta))^{-1})_{r,s} = \begin{cases} 1 & \text{if } r = s, \\ 0 & \text{if } r > s \text{ or if } r < s \text{ and } s \notin B, \\ 2^{p(B,r,s)} + t_{B,r,s}(\beta) & \text{if } r < s \text{ and } s \in B, \end{cases}$$

where $\sigma_r = -1$ if $r \in B$ and $\sigma_r = 1$ if $r \notin B$, $p(B, r, s) = |\{j \in B : r < j < s\}|$ and $t_{B,r,s}(\beta)$ is a polynomial in variables $\beta_{i,j}$ for $(i, j) \in \{(i, j) \in [n] \times B : i < j, (j < s) \text{ or } (j = s \text{ and } i \geq r)\}$ with no constant term. Moreover, $\beta_{r,s}$ appears in $t_{B,r,s}$ only in the linear term $\beta_{r,s}$ (that is, with coefficient 1).

From now on, we will write

$$(i, j) \prec (i', j') \quad \text{for} \quad (j < j') \text{ or } (j = j' \text{ and } i > i').$$

Note that \prec is a (strict) total ordering on $\{(i, j) \in [n]^2 : i < j\}$.

Let $B \subseteq [n]$ be a basis such that, for $m = \max B$, we have $i \equiv m + 1 \pmod{2}$ for each $i \in B \setminus \{m\}$. Then $q_m \cdot q_i = -1$ for each $i \in B \setminus \{m\}$, and hence

$$\begin{aligned} ((A_B(\beta))^{-1}q)_r &= (-1)^m \left(t_{B,r,m}(\beta) - \sum_{s \in B \setminus \{m\}} t_{B,r,s}(\beta) \right) \\ &= (-1)^m (\beta_{r,m} - t'_{B,r,m}(\beta)) \end{aligned} \tag{6}$$

for all $r < m$ such that $r \equiv m + 1 \pmod{2}$. Here $t'_{B,r,m}(\beta)$ is some polynomial in variables $\beta_{i,j}$ for $(i, j) \in \{(i, j) \in [n] \times B : i < j, (i, j) \prec (r, m)\}$ with no constant term.

3.6 Lemma. *Whenever $t'_{B,r,m}$ is defined, let*

$$\bar{\beta} = \max\{|\beta_{i,j}| : (i, j) \in [n] \times B, i < j, (i, j) \prec (r, m)\}.$$

If $\bar{\beta} < 1$, then $|t'_{B,r,m}(\beta)| < 9^{m-r+1}\bar{\beta}$.

Proof. The bound is actually very rough. First, by induction one can easily prove that the number of terms in each $t_{B,r,s}$ is at most 3^{s-r} , and the maximum coefficient in each $t_{B,r,s}$ is also at most 3^{s-r} . Since $\bar{\beta} < 1$, higher-degree terms can also be upper-bounded by $\bar{\beta}$. Hence $|t'_{B,r,m}(\beta)| < \sum_{s=r}^m 3^{2(s-r)}\bar{\beta} < 9^{m-r+1}\bar{\beta}$. \square

3.7 Lemma. *Let $r, m \in [n]$, $r < m$, $r \equiv m + 1 \pmod{2}$ and let $B, B' \subseteq [n]$ be bases such that $\max B = \max B' = m$ and that $i \equiv m + 1 \pmod{2}$ for all $i \in (B \cup B') \setminus \{m\}$. Then the polynomial $t'_{B,r,m}(\beta) - t'_{B',r,m}(\beta)$ is identically zero if and only if $B = B'$.*

Proof. Assume that $B \neq B'$. Without loss of generality, there exists some $s \in B \setminus B'$. Let $\tilde{t}_{B,r,m}(\beta_{r,s}), \tilde{t}_{B',r,m}(\beta_{r,s})$ be the univariate polynomials obtained from $t'_{B,r,m}(\beta), t'_{B',r,m}(\beta)$ respectively, by setting $\beta_{i,j} = 0$ for all $(i, j) \neq (r, s)$. Then $\tilde{t}_{B',r,m}(\beta_{r,s})$ is identically zero but $\tilde{t}_{B,r,m}(\beta_{r,s})$ is not. Hence $t'_{B,r,m}(\beta) - t'_{B',r,m}(\beta)$ is not identically zero. The converse implication is trivial. \square

Now let $r, m \in [n]$, $r < m$, $r \equiv m + 1 \pmod{2}$. Let

$$C = \{i \in [n] : r < i < m, i \equiv m + 1 \pmod{2}\}$$

and let

$$V' = \{(0 \oplus m) \oplus I : I \subseteq C\}.$$

Note that $|C| = (m - r - 1)/2$ and so $|V'| = 2^{(m-r-1)/2}$.

Furthermore, suppose that the values of $\beta_{i,j}$ are fixed for all $j < m$ and so are the values of $\beta_{i,m}$ for $i > r$, and that these values satisfy:

$$v, v' \in V', v \neq v' \implies t'_{B(v),r,m}(\beta) \neq t'_{B(v'),r,m}(\beta). \quad (7)$$

For each $v \in V'$, the direction of the edge between v and $v \oplus r$ in the USO induced by $\text{LCP}(M(\beta), q)$ is by (6) determined by the sign of the difference $\beta_{r,m} - t'_{B(v),r,m}(\beta)$. By (7), there are $|V'| + 1$ choices for $\beta_{r,m}$ so that the resulting USOs will differ from one another in the orientation of at least one of these edges. Moreover, by Lemma 3.7, the choices can be made in such a way that (7) is satisfied by the successor of (r, m) with respect to \prec (which is either $(r - 1, m)$, or $(m, m + 1)$).

The options to choose $\beta_{r,m}$ are, of course, not independent of the values of the other $\beta_{i,j}$'s. However, they depend only on the $\beta_{i,j}$'s with $(i, j) \prec (r, m)$. Hence it is possible to make the choices sequentially in the order given by \prec ; starting with $\beta_{1,2}$ and finishing with $\beta_{1,n}$.

Therefore the number of distinct USOs induced by $\text{LCP}(M(\beta), q)$ for various values of $\beta_{i,j}$, as described above, is at least

$$\prod_{m=1}^n \prod_{\substack{1 \leq r < m \\ r \equiv m+1 \pmod{2}}} \left(2^{(m-r-1)/2} + 1\right) = \prod_{m=1}^n \prod_{i=0}^{\lfloor m/2 \rfloor - 1} (2^i + 1) = 2^{\Omega(n^3)}.$$

Finally, it follows from Lemma 3.6 that the values of all $\beta_{i,j}$'s can be chosen to satisfy $|\beta_{i,j}| < 1$, so that $M(\beta)$ is a K-matrix. \square

3.4 The number of USOs from a fixed matrix

In this section, we prove the following

3.8 Theorem. *For a P -matrix $M \in \mathbb{R}^{n \times n}$, let $u(M)$ be the number of USOs determined by LCPs of the form $\text{LCP}(M, q)$ for $q \in \mathbb{R}^n$. Furthermore, define $u(n) = \max_M u(M)$, where the maximum is over all $n \times n$ P -matrices. Then*

$$u(n) = 2^{\Theta(n^2)}.$$

Proof. Let us first show the upper bound. For a fixed M , we consider the $n2^n$ hyperplanes of the form

$$\{x \in \mathbb{R}^n : (A_{B(v)}^{-1}x)_i = 0\}.$$

These hyperplanes determine an *arrangement* that subdivides \mathbb{R}^n into faces of various dimensions. Each face is an inclusion-maximal region over which the sign vector $(\text{sgn}(A_{B(v)}^{-1}x)_i) : v \in \{0, 1\}^n, i \in [n])$ is constant. The faces of dimension n are called *cells*; within a cell, the sign vector is nonzero everywhere. From Section 3.1 we know that $\text{LCP}(M, q)$ yields a USO whenever q is in some cell, and for all q within the same cell, $\text{LCP}(M, q)$ yields the same USO. Thus, the number of cells in the arrangement is an upper bound for the number $u(M)$ of different USOs induced by M . It is well-known [5] that the number of cells in an arrangement of N hyperplanes in dimension n is $O(N^n)$. In our case, we have $N = n2^n$ which shows that $u(M) = O((n2^n)^n) = 2^{O(n^2)}$ for all M .

For the lower bound, recall that we have constructed in Section 3.3 a K -matrix $M' \in \mathbb{R}^{(n-1) \times (n-1)}$ (resulting from fixing $\beta_{i,j}$ for all $j < n$), with the following property: for a suitable right-hand side q , $\text{LCP}(M, q)$ with

$$M = \begin{pmatrix} M' & b \\ 0 & 1 \end{pmatrix}$$

yields $2^{\Omega(n^2)}$ many different USOs in the subcube F corresponding to vertices with $v_n = 1$, when b is varied.

Since the subcube F corresponds to the solutions of $w - Mz = q$ that satisfy $w_n = 0$, we have $z_n = q_n$ within F . With $w' = (w_1, \dots, w_{n-1})^T$, $z' = (z_1, \dots, z_{n-1})^T$ and $q' = (q_1, \dots, q_{n-1})^T$, it follows that

$$w - Mz = q, \quad w^T z = 0, \quad w_n = 0$$

if and only if

$$w' - M'z' = q' - bq_n, \quad w'^T z' = 0, \quad z_n = q_n.$$

This is easily seen to imply that the induced USO in the subcube F is generated by $\text{LCP}(M', q' - bq_n)$. Thus, $u(M') = 2^{\Omega(n^2)}$, and the theorem is proved. \square

4 Locally uniform USOs and K-matrices

Finally we present a note on the relationship between K-matrices and locally uniform USOs.

4.1 Theorem. *Let M be a P-matrix. M is a K-matrix if and only if for all nondegenerate q , the USO induced by $\text{LCP}(M, q)$ is locally uniform.*

Proof. The “only-if” direction is Proposition 5.3 in [7]. For the if-direction, suppose that M is not a K-matrix. We will construct a vector q such that the induced USO violates (4). First, since M is not a K-matrix, there exists an off-diagonal entry $m_{ij} > 0$, $i \neq j$. W.l.o.g. assume that $\{i, j\} = \{1, 2\}$ and define

$$Q = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix}.$$

Let us now consider $B = \{1, 2\}$. Then

$$A_B = \left(\begin{array}{c|c} -Q & 0 \\ \hline 0 & I_{n-2} \end{array} \right),$$

and

$$A_B^{-1} = \left(\begin{array}{c|c} -Q^{-1} & 0 \\ \hline 0 & I_{n-2} \end{array} \right),$$

where

$$-Q^{-1} = \frac{1}{m_{11}m_{22} - m_{21}m_{12}} \begin{pmatrix} -m_{22} & m_{12} \\ m_{21} & -m_{11} \end{pmatrix}.$$

Since Q is a P-matrix, its determinant is positive, hence $-Q^{-1}$ has some positive off-diagonal entry. Suppose first that $m_{12} > 0$. Then we set $q = (-m_{12}, -(m_{22} + 1), 0, \dots, 0)$ and observe that

$$(A_B^{-1}q)_1 = \frac{-m_{12}}{m_{11}m_{22} - m_{21}m_{12}} < 0.$$

Slightly perturbing q such that it becomes nondegenerate will not change this strict inequality. But this is a contradiction to (4): at $B = \emptyset$, the edges in directions 1 and 2 are outgoing due to $q_1, q_2 < 0$, but at $B = \{1, 2\}$, the edge in direction 1 is *not* incoming as required by (4). If $m_{21} > 0$, the vector $q = (-(m_{11} + 1), -m_{21}, 0, \dots, 0)$ leads to the same contradiction. \square

Remark. It may be more interesting to answer the following open question: Is it true that every locally uniform P-USO is a K-USO?

Acknowledgments

The third author is supported by the project ‘A Fresh Look at the Complexity of Pivoting in Linear Complementarity’ no. 200021-124752 / 1 of the Swiss National Science Foundation.

References

- [1] R. Chandrasekaran. A special case of the complementary pivot problem. *Opsearch*, 7:263–268, 1970.
- [2] S.-J. Chung. NP-completeness of the linear complementarity problem. *J. Optim. Theory Appl.*, 60(3):393–399, 1989.
- [3] R. W. Cottle, J.-S. Pang, and R. E. Stone. *The Linear Complementarity Problem*. Computer science and scientific computing. Academic Press, 1992.
- [4] M. Develin. LP-orientations of cubes and crosspolytopes. *Adv. Geom.*, 4(4):459–468, 2004.
- [5] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Heidelberg, West Germany, 1987.
- [6] K. Fischer and B. Gärtner. The smallest enclosing ball of balls: combinatorial structure and algorithms. *Internat. J. Comput. Geom. Appl.*, 14(4–5):341–378, 2004.
- [7] J. Foniok, K. Fukuda, B. Gärtner, and H.-J. Lüthi. Pivoting in linear complementarity: Two polynomial-time cases. *Discrete Comput. Geom.*, 42(2):187–205, 2009.
- [8] B. Gärtner. Randomized algorithms – An introduction through unique sink orientations. Lecture Notes, ETH Zurich, 2004.
- [9] B. Gärtner, W. D. Morris, Jr., and L. Rüst. Unique sink orientations of grids. *Algorithmica*, 51(2):200–235, 2008.
- [10] B. Gärtner and I. Schur. Linear programming and unique sink orientations. In *Proceedings of the 17th annual ACM-SIAM symposium on Discrete algorithms (SODA’06)*, pages 749–757, Miami, Florida, 2006.
- [11] L. Klaus. On classes of unique-sink orientations arising from pivoting in linear complementarity. Master’s thesis, ETH Zurich, 2008.
- [12] O. L. Mangasarian. Linear complementarity problems solvable by a single linear program. *Math. Programming*, 10(2):263–270, 1976.
- [13] J. Matoušek. The number of unique-sink orientations of the hypercube. *Combinatorica*, 26(1):91–99, 2006.
- [14] N. Megiddo. A note on the complexity of P-matrix LCP and computing an equilibrium. RJ 6439, IBM Research, Almaden Research Center, 650 Harry Road, San Jose, California, 1988.

- [15] H. Miyazawa. Cube orientations and properties. Project Report, ETH Zurich, 2001.
- [16] W. D. Morris, Jr. Distinguishing cube orientations arising from linear programs. Manuscript, 2002.
- [17] W. D. Morris, Jr. Randomized pivot algorithms for P-matrix linear complementarity problems. *Math. Program., Ser. A*, 92(2):285–296, 2002.
- [18] W. D. Morris, Jr. and J. Lawrence. Geometric properties of hidden Minkowski matrices. *SIAM J. Matrix Anal. Appl.*, 10(2):229–232, 1989.
- [19] K. G. Murty. *Linear Complementarity, Linear and Nonlinear Programming*, volume 3 of *Sigma Series in Applied Mathematics*. Heldermann, Berlin, 1988.
- [20] J.-S. Pang and R. Chandrasekaran. Linear complementarity problems solvable by a polynomially bounded pivoting algorithm. *Math. Programming Stud.*, 25:13–27, 1985.
- [21] R. Saigal. A note on a special linear complementarity problem. *Opsearch*, 7:175–183, 1970.
- [22] I. Schurr and T. Szabó. Finding the sink takes some time: An almost quadratic lower bound for finding the sink of unique sink oriented cubes. *Discrete Comput. Geom.*, 31(4):627–642, 2004.
- [23] I. A. Schurr. *Unique Sink Orientations of Cubes*. PhD thesis, ETH, Zürich, 2004.
- [24] A. Stickney and L. Watson. Digraph models of Bard-type algorithms for the linear complementarity problem. *Math. Oper. Res.*, 3(4):322–333, 1978.
- [25] T. Szabó and E. Welzl. Unique sink orientations of cubes. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS'01)*, pages 547–555, 2001.
- [26] H. E. Warren. Lower bounds for approximation by nonlinear manifolds. *Trans. Amer. Math. Soc.*, 133:167–178, 1968.